

Learning to Bid

– An Experimental Study of Bid Function Adjustments in Auctions and Fair Division Games

Werner Güth, Radosveta Ivanova,
Manfred Königstein, Martin Strobel
Humboldt University at Berlin[⌘]

August 25, 1999

Abstract

We examine learning behavior in auctions and fair division games with independent private values under two different price rules, first and second price. Participants face these four games repeatedly and submit complete bid functions rather than single bids. This allows us to examine whether learning is influenced by the structural differences between games. We find that within the time horizon which we investigate, learning does not drive toward risk neutral equilibrium bidding and characterize some features of observed learning: Bid functions are adjusted globally rather than locally, decision time matches the sequencing structure of game types, game rules do matter, and directional learning theory offers a partial explanation for bid adjustments. The evidence supports a cognitive approach to learning.

JEL classification: D44, C91

Keywords: Auctions, Fair Division Games, Learning

[⌘]Institut für Wirtschaftstheorie III, Humboldt-Universität zu Berlin, Spandauer Str. 1, 10178 Berlin, email: {gueth,ivanova,mkoenig,strobel}@wiwi.hu-berlin.de

1. Introduction

Recently, learning has become a major topic in economic research. Whereas in former times one either relied on (common knowledge of) rationality or that markets will drive out irrational modes of behavior, one now is interested in the processes of behavioral adjustments and open to what may be their results. In our view, this interest in learning is most welcome and inspires behavioral economics, especially since learning models explicitly account for decision makers imperfect abilities to collect, store, retrieve, and analyze information.

One way to structure the different approaches to learning which are applied in theoretical as well as empirical research is to separate cognitive- from non-cognitive learning theories (see Selten, 1997, who distinguishes between 'cognition' as reasoning and 'adaptation' as routine adjustment without reasoning). Among the cognitive models are e.g. 'best-reply dynamics', 'imitation' and 'directional learning' (see, for instance, Selten and Buchta, 1998). Among non-cognitive models 'reinforcement learning' (see Bush and Mosteller, 1955 and Roth and Erev, 1995), 'imitation' (Vega-Redondo, 1997), and 'replicator dynamics' (Taylor and Jonker, 1978) are quite well-known.

While some cognitive-learning theories may be criticized of not going far enough in modelling agents' intellectual restrictions — for instance, best-reply dynamics still assume that economic agents are perfect maximizers — the non-cognitive ones might be accused of going too far by abandoning basically every reasoning. We think that behavioral adjustment should be explained by some sort of boundedly rational reasoning, allowing for adjusting behavior in response to previous experiences, both by reevaluation of parameters and by cognitive adaptations, e.g. in the sense of including new variables into the reasoning process. Our study presents evidence in support of this view. For an investigation that explicitly separates the respective predictions of cognitive and non-cognitive adaptation see Abbink, Bolton, Sadrieh, and Tang (1996).

In our experiment we asked subjects to develop bid functions¹ for four different types of games, the first (respectively second) price auction and the first (respectively second) price fair division game. A bid function specifies a bid for each of 11 possible private values. After submitting bid functions a value is drawn randomly and independently for each subject and the game outcome is determined accordingly. Subjects participated in 36 games against randomly matched opponents

¹For other experiments in which subjects had to submit bid functions see Selten and Buchta (1998) as well as Güth (1998).

and with the game type changing after every three rounds of bidding.

Developing a bid function is more demanding than deciding about a bid for a single value; which is the usual practice in studies on bidding behavior (see e.g. Kagel, 1995). Furthermore, since we exposed subjects to different games, the decision task can be considered as rather complex. Nonetheless, in this environment we found structures of learning behavior which we interpret as cognitive though boundedly rational. While many studies on learning focus on rather simple games and an extremely long time horizon we think it is important to look also at learning in more complex games with a short or intermediate time horizon. We will discuss this in more detail later on.

As indicated within this paper we focus on learning aspects of bidding behavior. In a companion paper (Güth, Ivanova, Königstein and Strobel, 1999) we show that in most cases the bid functions are increasing, linear (or almost linear) and in line with the comparative statics of equilibrium bidding for the different game types (even though, in general, behavior does not match risk neutral equilibrium behavior). Furthermore, in that paper we compare efficiency aspects and price aspects of the various game types.

Our paper here is structured as follows: Section 2 describes the experimental games and their theoretical benchmark solutions. Furthermore we inform about the experimental procedures we applied and the payments to subjects. Section 3 presents our empirical results and section 4 concludes.

2. Auctions and Fair Division Games

2.1. Games and Theoretical Solutions

Bidding behavior is a favorite topic in experimental economics (see for a selective survey Kagel, 1995). Like in auction theory one distinguishes sequential and sealed-bid auctions. We will focus here on sealed bid-experiments in which a single object is to be allocated and for which each potential buyer has an independent private value. We investigate four different allocation rules which we refer to as game types (see table 2.1): First Price Auction (A1), Second Price Auction (A2), First Price Fair Division Game (F1) and Second Price Fair Division Game (F2). Fair division games differ from auctions in that the price at which the object is sold gets distributed among all bidders. In auctions the price is earned by an outside agent, the seller. While the use of auctions to solve allocation problems is common, fair division games may be less familiar. For instance, allocating inheritance is a

Price Rule	Auction	Fair Division Game
price = highest bid	A1	F1
price = 2nd highest bid	A2	F2

Table 2.1: The four game types.

Price	Auction	Fair Division Game
highest bid	$b_i^a(v_i) = \frac{n_i-1}{n}v_i$ $E(p^a) = \frac{n_i-1}{n+1}$ $E(\mathcal{Y}_i^a(v_i)) = \frac{v_i^n}{n}$	$b_i^a(v_i) = \frac{n}{n+1}v_i$ ² $E(p^a) = \frac{n}{n+1}$ ² $E(\mathcal{Y}_i^a(v_i)) = \frac{v_i^n}{n} + \frac{n_i-1}{n(n+1)}$
2nd highest bid	$b_i^a(v_i) = v_i$ $E(p^a) = \frac{n_i-1}{n+1}$ $E(\mathcal{Y}_i^a(v_i)) = \frac{v_i^n}{n}$	$b_i^a(v_i) = \frac{n}{n+1}v_i + \frac{1}{n+1}$ $E(p^a) = \frac{n^2+1}{(n+1)^2}$ $E(\mathcal{Y}_i^a(v_i)) = \frac{v_i^n}{n} + \frac{n_i-1}{n(n+1)}$

Table 2.2: Bidding function, expected price and expected payoff according to the risk neutral equilibrium for the four game types.

real life situation which resembles a fair division game. The object is collectively owned by the heirs who, in many cases, are the only bidders. Similar problems result when a joint venture is terminated.²

In both, auctions as well as fair division games, the 'axiom of envy-free net trades with respect to bids' (Güth, 1986) implies that the object should be allocated to the highest bidder at a price which does not exceed the highest bid and which does not fall below the second highest bid. The two price rules we investigate, highest respectively 2nd highest price, are the two polar cases satisfying this axiom. Besides, comparing bidding behavior under these two price rules is familiar in experimental studies.³

Let v_i be a bidder's private value for the object to be sold, and suppose v_i is drawn for each player $i = 1, \dots, n$ independently from a uniform distribution on the

²For an experimental study on a related topic, the so called zero-revenue auctions, see Franciosi, Isaac, Pingry, and Reynolds (1993). However, since these are multi-unit auctions we do not compare their results to ours.

³See the survey by Kagel (1995).

unit interval. If all bidders are risk neutral, the equilibrium bid function $b_i^*(v_i)$; expected equilibrium price $E(p^*)$ and expected equilibrium payoff $E(\pi_i^*(v_i))$ are as shown in table 2.2. For a derivation of these results see Güth and van Damme (1986).

2.2. Experimental Games and Procedures

In our experiment the private values v_i did not vary continuously, but were drawn from the set

$$V = \{50; 60; 70; 80; 90; 100; 110; 120; 130; 140; 150\}$$

with all values $v_i \in V$ being equally likely. These values are denoted in an artificial currency ECU (experimental currency unit) at which subjects could resell the object to the experimenter. For ease of comparison of the empirical bids b_i and values v_i with the theoretical solution given above all our analysis will be done for normalized bids \tilde{b}_i and values \tilde{v}_i :

$$\tilde{v}_i = \frac{v_i - 50}{100}$$

$$\tilde{b}_i = \frac{b_i - 50}{100} :$$

Accordingly the space of possible values is $V = \{0; 0.1; \dots; 1\}$. When we refer to the theoretical benchmark case as described above, we essentially neglect the discreteness of V :

Within a session each subject participated in 36 consecutive games of the four different types. Nine subjects formed a session group. In each of the 36 periods they were randomly partitioned into three groups of three bidders. The number of bidders involved in each game ($n = 3$) was commonly known, but not their identity. All subjects in all sessions played the same sequence of games. Within periods $t = 1$ to 3 they played A1, within $t = 4$ to 6 they played A2, in $t = 7$ to 9 the game type was F2 and in $t = 9$ to 12 it was F1. This comprised the first block of 12 games. Then they played block 2 (periods 13 to 24) and 3 (periods 25 to 36) in the same sequence as block 1.

Most participants were students of economics or business administration of Humboldt-University. They had been invited by leaflets to participate in an experiment announced to last about three hours, and sessions actually took about that long. After entering the laboratory they were placed at isolated computer

terminals. Communication among participants was not allowed during the session. While reading the instructions (see the appendix) they could privately ask for clarifying questions or for help regarding handling the PC.

In each game they had to submit a complete bidding strategy (bid vector) $b_i(v_i)$. Thus, they had to enter a bid for each of the 11 values $v_i \in V$: The actual value v_i^0 was drawn thereafter. Payments were determined according to the game rules and using the submitted bidding strategies.⁴ Subjects were informed on screen about v_i^0 , about whether or not they were buyer, about the price p at which the object was sold and about their own payoff π_i in that game. Then the next game followed. Appendix B shows some sample screen shots.

Each game type applied nine times. In the ...rst of these nine games the bid screen was blank and each subject had to enter a new vector of 11 bids (one for each $v_i \in V$). In later periods the last bid vector for the same game type was displayed as default. It could be revised or submitted as it is. Of course, this may favor the status quo and may work against adjusting behavior over time. We did it for practical reasons. If subjects do not want to always adjust all bids, this saves time and helps to prevent them from getting bored by the task. Altogether we ran 6 sessions and collected 1944 bidding strategies (54 subjects times 36 games).

2.3. Payments

Subjects total earnings out of the 36 games ranged between 31 DM and 96 DM with a mean of 56 DM (about 33 US\$ at the time of the experiment including a show up fee of 10 DM). In the ...rst half of the sessions we used the same conversion rate for ECU (Experimental Currency Unit) into cash: 1 ECU = 0.05 DM. Theoretically and practically this generates rather asymmetric monetary incentives for auctions versus fair division games. Güth (1998) tried to guarantee equal monetary incentives by adjusting the conversion rate such that equilibrium profits were equal for $v_i = 0:5$. Instead we used actually observed profits of the ...rst three sessions and adjusted the conversion rate to induce equal expected payoffs based on observed behavior. This meant for sessions 4 to 6 that one ECU was worth DM 0.2857 in auctions and DM 0.02857

Essentially this means that we had a payoff-treatment: 3 sessions with equal conversion rate and 3 sessions with unequal conversion rate. Theoretically these

⁴The strategy method obviously provides more information than collecting only one bid for a single value. But since ex-post only one component of the bid vector is payoff-relevant, it lowers the incentives of bidding at each single value. By restricting the set V we have tried to achieve a reasonable compromise.

payoff differences are irrelevant. And, since in all data analyses we ran, we did not find them being relevant, we will not discuss them any further.

3. Results

3.1. Monotonicity and Convergence of Bid Adjustments

Before investigating whether subjects learn any specific kind of behavior, one might ask whether they did learn something at all. If subjects do learn something, one should observe — after sufficient time — some stationary behavioral pattern. Here, this would mean that individuals adjust their bid functions and with experience approach some stable individual bid function.

Each subject i played each type of game nine times. For every type we will refer to the bid function in the ninth play as i 's final bid function. Hence the final bid functions of game type A1, A2, F2, and F1 are $b_{i;27}(v_i)$, $b_{i;30}(v_i)$, $b_{i;33}(v_i)$, and $b_{i;36}(v_i)$, respectively. To measure bid function adjustments for game type A1 we calculated separately for each individual the Euclidean distance $D_i^{A1}(t) = \|b_{it}(v_i) - b_{i;27}(v_i)\|$ between i 's bid function $b_{it}(v_i)$ in period t and i 's final bid function where $t \in \{1; 2; 3; 13; 14; 15; 25; 26\}$, i.e. t is a period in which game type A1 was played. Analogously we calculated $D_i^{A2}(t)$, $D_i^{F2}(t)$, and $D_i^{F1}(t)$. We consider an adjustment process as monotone, if $D_i^j(t)$ with $j \in \{A1, A2, F2, F1\}$ is decreasing in t . Furthermore, a monotone adjustment process will be called 'convergent', if $D_i^j(t)$ decreases more rapidly in earlier than in later periods, i.e. if $D_i^j(t)$ is convex. Both, monotone processes and (even more) converging processes, will be interpreted as evidence for learning.

For classification of the observed processes we used slightly weaker criteria than those described above in order to allow for some error. Specifically, we fitted a piecewise-linear regression line to the data with $D_i^j(t)$ as dependent and t as independent variable, and allowing for a kink of this line after 4 (out of the 8) periods. Accordingly, a process is regarded monotone, if both slope coefficients are negative. If in addition the coefficient is smaller in absolute value for later periods, the process is convergent.

Table 3.1 displays the relative frequencies of observed monotone, respectively convergent adjustment processes. As one can see, between 61% and 74% of all individual bid functions exhibit a convergent adjustment process; between 92% and 100% are monotone. Thus, even though we do not know yet what specific kind of bidding behavior subjects learn, we know that they do learn something.

Adjustment Process	Game Type			
	A1	A2	F2	F1
Monotone	92%	95%	96%	100%
Monotone and Convergent	61%	69%	74%	70%
Monotone and Not Convergent	31%	26%	22%	30%
Not Monotone	8%	5%	4%	0%
Sum (mon. and non-mon.)	100%	100%	100%	100%

Table 3.1: Classification of adjustment processes for different game types.

3.2. No Learning of Risk Neutral Equilibrium Bidding

Many theoretical models in auction theory work with the assumption that bidders are risk neutral (see e.g. the survey by Wolfstetter, 1996). It was shown in Güth et al. (1999) for the same data set analyzed here that by and large the risk neutral equilibrium (RNE) fits the data poorly. Nevertheless, since there is a learning process going on, it could well be that this process works into the direction of the RNE. In this case one might speculate whether subjects will play according to the RNE if they were given more time to gain experience and to learn. However, figures 3.1-3.4 suggest that this is not the case. They present time paths of estimates of the aggregate bid functions of each session. Specifically, based on piecewise-linear regressions of the aggregate bid function for each session, they show the predicted bids for $v_i = 0$; $v_i = 0.5$; and $v_i = 1$ for the four game types.⁵ A reference line in each figure indicates the RNE bid. We hardly observe any movement towards it. So, we conclude that the learning process does not drive towards the RNE.

⁵To give the RNE a better chance we excluded some data which were obviously problematic. Remember that altogether we collected 1944 individual bid functions. Güth et al. (1998) showed that 98% of these were strictly increasing and that most of them were quite accurately predicted (according to the coefficient of determination R^2) by a piecewise-linear model — i.e., $b_i(v_i)$ is piecewise-linear in v_i — allowing for a kink of the regression line at $v_i = 0.5$: In computing the estimates of the aggregate session bid function, which we present here, we therefore excluded individual bid functions which were not strictly increasing, and furthermore those with an R^2 of the piecewise-linear model smaller than 80%. Taken together we thereby excluded 84 (4%) of the bid functions.






Figure 3.1: Time paths of estimated bids in ...rst price auctions (A1)




Figure 3.2: Time paths of estimated bids in second price auctions (A2)




Figure 3.3: Time paths of estimated bids in second price fair division games (F2)




Figure 3.4: Time paths of estimated bids in ...rst price fair division games (F1)

3.3. Cognitive versus Non-Cognitive Learning

It has become quite popular to explain economic behavior as being driven by some non-cognitive dynamic process. Examples for theoretical work in this field are models in evolutionary game theory (see e.g. Weibull, 1995). An example for empirical work trying to explain laboratory behavior by reinforcement learning is Roth and Erev (1995). Most of these studies rely on rather simple game models⁶ and an extremely long time horizon. Although such experiments are obviously needed, their results may not be applicable to many real life-decision problems. One usually does not play 500 or 1000 times. So, a short- or intermediate- rather than a long term perspective may be more important.

We conjecture that by a cognitive short- or mid-term learning process the boundedly rational economic agent eliminates some bad decisions to reach a stable behavioral pattern at which some aspiration level is reached. After reaching this stable pattern — this does not necessarily mean constant play — the agent might occasionally try some new strategy to see whether he may reach an even better outcome. At this stage behavior might resemble 'trial and error'-experimentation or some other non-cognitive dynamics. If a game is rather simple, a stable behavioral pattern may be developed from the start, and all subsequent variation in decisions might follow non-cognitive dynamics. So, while we do not deny that some economic behavior can be captured by a non-cognitive dynamic process, we believe that the largest adjustments in behavior take place within a short or intermediate time span and are based on a cognitive, though boundedly rational, analysis of the situation. However, whether one perspective on learning is superior to the other has to be found out via empirical research. Following this programmatic discussion, the next sections will provide evidence which seems to support a cognitive approach.

3.4. Local Adjustments versus Global Adjustments

While in each game each subject had to enter a complete bid function, only one private value v_i^0 was actually drawn (in each game). So, within that game only the bid submitted for this value $b_i(v_i^0)$ was payoff relevant. Accordingly, the informational feedback received by individual i after each game — i.e., i 's value v_i^0 ; the price p ; whether or not i bought the object and i 's profit — might suggest whether $b_i(v_i^0)$ should be adjusted in future periods, but it does not tell anything

⁶For instance, in order to apply their learning model upon the ultimatum game — which is not a too complicated game anyway — Roth and Erev simplify it even further.




Figure 3.5: Frequency distributions for the number of bid functions changes for the different game types.

regarding bids $b_i(v_i^{00})$ for all values $v_i^{00} \in v_i^0$. A naturally arising question is therefore whether bid functions were adjusted only 'locally' at v_i or rather 'globally' at all values.

Figure 3.5 shows frequency distributions for the number of bid functions changes for the different game types. Note that the maximal number of changes is 11; this means that a bid vector is changed in each component. We observe that if a bid function is changed at all, it is simultaneously changed at all values in most cases.⁷ We therefore conclude that bid functions are adjusted globally rather than locally.

⁷A single subject plays each game type nine times. So, it can change its bid function for each game in at most eight periods (change periods). The maximal number of change periods per subject is therefore 32 and the maximal number of change periods per game type is 432 (54 subjects times 8). The percentages of change periods (observed change periods divided by maximal change periods times 100) are 54% in A1, 48% in A2, 34% in F2 and 33% in F1.




Figure 3.6: Average decision time by periods 1 to 36.

3.5. Sequence Structure of Games and Decision Time

As explained above we chose a specific sequencing of playing the 4 types of games. Each game was played 3 times in a row before the next type of game was played for 3 periods. 12 periods (4 times 3 periods) formed 1 block. The sequence within each block was A1, A2, F2, F1. Altogether subjects played 3 blocks (36 games). While the game rules differed between the 4 game types, the format of subjects' decisions were the same. Figure 3.6 shows that the average decision time closely matches the sequencing structure of the games. Decision time is high at the beginning and decreases over periods, but it does not decrease too fast. If subjects followed a non-cognitive learning approach, we do not see why decisions should take less time in the end than in the middle of the sequence. Furthermore, Figure 3.6 shows that decision time 'jumps up' whenever the game type changes; i.e. at periods 1, 4, 7, 10, ...34. Within the first block (periods 1, 4, 7, and 10) this is natural, since all subjects had to type in a new bid vector; in later periods, however, they could rely on their former strategy for the same game type and just click the OK-button. But, the displayed path also shows spikes in periods 13, 16, ..., 34 where the above explanation does not hold. If one takes the time span subjects need to come up with a decision as an indicator of their cognitive effort,

this structure is quite plausible.⁸

3.6. Directional Learning

After each play a participant was informed about the realization of his own value v_i^0 , whether or not he became buyer, the price p at which the object was sold and his own profit. Clearly, such feedback allows for directional learning in the sense of ex post-comparisons of the chosen versus alternative bids (Selten and Buchta, 1998). In the following we want to examine 'Directional Learning'-theory within a bidding context (see Selten and Buchta, 1998, and Güth, 1998). This theory tries to predict the direction of strategy adaptation after information feedback: If another strategy would have yielded a better outcome according to the received information a participant is expected to move (i.e., adjust his strategy) into this direction.

So, in first price auctions a subject i who became buyer in period t_{j-1} should lower his bid in t , since a bid reduction might have increased his payoff in t_{j-1} : On the contrary, a non-buyer j whose value in t_{j-1} was above the price, $v_j^0 > p$; could have made a positive profit by a higher bid and therefore should increase his bid in t : If $v_j^0 < p$ non-buyer j could not possibly have made a positive profit. Hence learning direction theory does not make a prediction in this case (see Selten and Buchta, 1998).

By similar arguments we developed conditions for directional predictions for all four game types. They are summarized in table 3.2. We used a slightly different criterion in A1 than Selten and Buchta for the following reasons. Consider buyer i who earned a positive profit in t_{j-1} : Although a bid reduction might have increased his profit, he might as well have lost money by becoming a non-buyer. Selten and Buchta argue that 'it is not clear how high the second bid was but higher profit could have been made by some lower bid' (p. 12). But this holds only for continuous bids. Since in our experiment bids are discrete, even a small bid reduction may lead to i becoming a non-buyer. However, if buyer i earned a non-positive profit in t_{j-1} ; a bid reduction could have never reduced but may have increased his profit. This is the condition we applied in A1. Analogous reasoning results in the other conditions displayed in table 3.2.

Altogether we found 318 (14.18%) out of 1728 cases where directional learning is

⁸We compared the results for sessions with even conversion rates (ECU to German Mark) for auctions and fair division games to those with uneven conversion rate, to see whether decision time is increasing in payoff, but found no systematic difference.

Predicted Direction of Bid Change	Condition for Prediction			
	A1	A2	F2	F1
bid #	Buyer and $b_i > 0$	Buyer and $b_i < 0$	Buyer and $b_i < \frac{1}{3}p$	Buyer and $b_i > \frac{1}{3}p$
bid "	Not Buyer and $p < v_i^0$	Not Buyer and $p < v_i^0$	Not Buyer and $p < v_i^0$	Not Buyer and $p < v_i^0$
no prediction	Otherwise	Otherwise	Otherwise	Otherwise

Table 3.2: Conditions for applying directional learning theory and predicted directions of bid changes.

applicable. The numbers and relative frequencies of correct and wrong bid changes as well as the numbers and frequencies of no bid changes are displayed in table 3.3 (rows denoted by "Sum"). Accordingly, the percentage of wrong predictions is 5% or smaller for all four game types; on average it is 4%. The frequency of right predictions is on average 41%, while bids remain unchanged in 55% of the cases. Observe that the percentage of right predictions is substantially smaller in fair division games than in auctions. We attribute this to the more complex structure of fair division games in contrast to auctions. Therefore subjects face more difficulties to figure out the correct direction via cognitive thinking.

In the rows denoted by "Change period" and "Other period" we break down the information of the "Sum"-rows by different types of periods: periods in which the game type had just changed (change periods) and other periods. Note that in change periods the relevant feedback information to be applied by a 'directional learner' was received 10 periods before. In other periods it was received just before the current period. A natural implication of boundedly rational directional learning is therefore that the predictive success should be smaller for change periods than for other periods. Table 3.3 shows that this is indeed the case for all four game types, although the effect is partly marginal.

We sum up our findings regarding directional learning as follows:

1. Due to our criteria for predicting the direction of bid changes, about 18% of all observed bids can be considered as test cases for directional learning theory.
2. We find violations of the theory in less than 6% of the cases for each game

Game type	Observed direction				Sum (100%)
	Right	Wrong	No bid change		
A1 Sum	34 (56%)	2 (3%)	25 (41%)		61
Change period	5 (50%)	1 (10%)	4 (40%)		10
Other period	29 (57%)	1 (2%)	21 (41%)		51
A2 Sum	39 (57%)	1 (1%)	29 (42%)		69
Change period	11 (55%)	0 (0%)	9 (45%)		20
Other period	28 (57%)	1 (2%)	20 (41%)		49
F2 Sum	33 (34%)	5 (5%)	59 (61%)		97
Change period	7 (33%)	1 (5%)	13 (62%)		21
Other period	26 (34%)	4 (5%)	46 (61%)		76
F1 Sum	25 (27%)	4 (5%)	62 (68%)		91
Change period	4 (15%)	3 (11%)	20 (74%)		27
Other period	21 (33%)	1 (1%)	42 (66%)		64
All game types, Sum	131 (41%)	12 (4%)	175 (55%)		318
Change period	27 (35%)	5 (6%)	46 (59%)		78
Other period	104 (43%)	7 (3%)	129 (54%)		240

Table 3.3: Predictive success of directional learning in change periods versus other periods.

type. If the bids are changed at all (and directional learning is applicable), they are mostly changed into the right direction. The conditional frequency for adjusting into the right direction is 92% on average (which can be calculated immediately from table 3.3).

3. Learning direction theory predicts slightly better for recent feedback compared to more distant feedback.
4. Learning direction theory predicts better for game types which are less complex.

4. Conclusions

In (applied) game theory one has theoretically assumed that rationality of all strategically interacting agents is commonly known. Although one often justifies this by an "as if"-interpretation (irrational behavior will be eliminated by market forces), the new interest in learning and evolution has inspired many thorough studies when rationality can be expected and when not.

There are, of course, many ways to formulate the dynamics of learning and evolution. Best reply dynamics for instance, only questions the assumption that rationality is commonly known, but it does not question at all individual decision rationality. Other learning and evolutionary dynamics like reinforcement learning (Bush and Mosteller, 1955 and Roth and Erev, 1995) or replicator dynamics (Weibull, 1995) deny any cognition and view behavior as being driven by past results, especially past (reproductive) success.

In our study here we presented evidence for cognitive rather than non-cognitive adjustment of bidding behavior. While the data reject learning of RNE bidding, which is the benchmark solution for risk neutral and perfectly rational bidders, subjects do learn something. Almost all adjustment paths are monotone and most are convergent. We do not want to speculate about where this learning will eventually settle, but characterize some features of the process.

First, bid functions are adjusted globally rather than locally. So, subjects 'interpret' the local feedback as having informational content for other or even all components of the bid vector. A reasonable cognitive principle that might explain such behavior is 'generalization'. Non-cognitive approaches could possibly model this behavior as parameter learning, e.g. by allowing adjustments of a proportional or absolute degree of under-, respectively overbidding that is applied to all components of the bid vector.

Secondly, while overall the time subjects use for deciding is decreasing by periods, it closely matches the sequencing structure of the different games types. A plausible explanation is that decision time represents cognitive effort.

Last but not least, we have argued that some cognitive learning theory, directional learning, offers at least a partial explanation of the observed changes in bidding behavior.

Nonetheless, more work needs to be done until the cognitive processes of bidding adjustments in auctions and fair division games are well understood.

References

- [1] Abbink, K., G. Bolton, A. Sadrieh, and F. Tang (1996), Adaptive learning versus punishment in ultimatum bargaining, Discussion Paper, SFB-303-B-381, University of Bonn.
- [2] Bush, R. and F. Mosteller (1955), Stochastic Models for Learning, New York: Wiley.
- [3] Franciosi, R., R. M. Isaac, D. E. Pingry, S. S. Reynolds (1993), An Experimental Investigation of the Hahn-Noll Revenue Neutral Auction for Emissions Licences, *Journal of Environmental Economics and Management*, 24, 1-24
- [4] Güth, W. (1986), Auctions, public tenders, and fair division games: An axiomatic approach, *Mathematical Social Sciences*, 11, 283-294.
- [5] Güth, W. (1997), Boundedly rational decision emergence — A general perspective and some selective illustrations, Discussion Paper, SFB-373-32, Humboldt-University at Berlin.
- [6] Güth, W. (1998), On the effects of the pricing rule in auctions and fair division games - An experimental study, in: D. Budescu, I. Erev, R. Zwick (eds.), *Games and Human Behavior*, Mahwah (N.J.): Lawrence Erlbaum Ass.
- [7] Güth, W. and E. van Damme (1986), A comparison of pricing rules for auctions and fair division games, *Social Choice and Welfare*, 3, 177 - 198.

- [8] Güth, W., R. Ivanova, M. Königstein and M. Strobel (1999), Auctions and fair division games under different price rules: Individual bid functions, prices, and efficiency rates, Working Paper, Humboldt University at Berlin.
- [9] Kagel, J. (1995), Auctions: A Survey of Experimental Research, in: J. Kagel and A. R. Roth (eds.), *The Handbook of Experimental Economics*, Princeton (N.J.): Princeton University Press.
- [10] Roth, A. E. and I. Erev (1995), Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term, *Games and Economic Behavior* 8, 164 - 212.
- [11] Selten, R. (1997), Features of experimentally observed bounded rationality, Discussion Paper, SFB-303-B-421, University of Bonn.
- [12] Selten, R. and J. Buchta (1998), Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions, in: D. Budescu, I. Erev, R. Zwick (eds.), *Games and Human Behavior*, Mahwah (N.J.): Lawrence Erlbaum Ass.
- [13] Taylor, P. and L. Jonker (1978), Evolutionary stable strategies and game dynamics, *Mathematical Biosciences*, 40, 145-156.
- [14] Vega-Redondo, F. (1997), The Evolution of Walrasian Behavior, *Econometrica*, 65, 375-384.
- [15] Weibull, J. (1995): *Evolutionary Game Theory*, Cambridge (M.A.): The MIT Press.
- [16] Wolfstetter, E. (1996), Auctions: An Introduction, *Journal of Economic Surveys*, Vol. 10 (4), 367-420.

A. Instructions⁹

Please read these instructions carefully. They are identical for all participants.

During the experiment you will take part in several auctions. In every auction a fictitious commodity is for sale which you can resell to the experimenters. You

⁹This is a shortened and translated version of the instructions. For the original instructions (in German), please contact one of the authors.

are one of three bidders. Each bidder has his own private reselling value v which can be 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, or 150 ECU (Experimental Currency Unit) and is independently drawn. Each value appears with the same probability.

Before you get to know your individual reselling value v you have to place a bid for every possible v :

$b(50); b(60); b(70); \dots; b(150)$.

After every bidder in your group has placed his bid vector your actual bid is determined by $b(v)$. The bidder with the highest bid buys the commodity and pays a price according to the pricing rule. Then he sells the commodity to the experimenter and receives his reselling value. The other bidders do not pay anything and do not receive the commodity. If there are two or three highest bids, the buyer is chosen by random.

There are 4 different types of auctions. In type 1 and 2 the auction revenue is kept back by the experimenter whereas in type 3 and 4 the auction revenue is equally divided among the bidders. In auction types 1 and 4 the price corresponds to the highest actual bid. In auction types 2 and 3 the price which has to be paid corresponds to the second highest actual bid.

Type 1 (First Price Auction)

- ² Price = highest bid ($p = b_1$)
- ² Bidder with highest bid becomes buyer. He pays p .
- ² Revenue (p) is kept back by the experimenter.
- ² Profit of buyer: $v_i - p = v_i - b_1$
- ² Profit of non-buyers: 0

Type 2 (Second Price Auction)

- ² Price = second highest bid ($p = b_2$)
- ² Bidder with highest bid becomes buyer. He pays p .
- ² Revenue (p) is kept back by the experimenter.

² Profit of buyer: $v_i - p = v_i - b_2$

² Profit of non-buyers: 0

Type 3 (Second Price Fair Division Game)

² Price = second highest bid ($p = b_2$)

² Bidder with highest bid becomes buyer. He pays p .

² Revenue (p) is distributed among the bidders.

² Profit of buyer: $v_i - p + \frac{1}{3}p = v_i - b_2 + \frac{1}{3}b_2 = v_i - \frac{2}{3}b_2$

² Profit of non-buyers: $\frac{1}{3}p = \frac{1}{3}b_2$

Type 4 (First Price Fair Division Game)

² Price = highest bid ($p = b_1$)

² Bidder with highest bid becomes buyer. He pays p .

² Revenue (p) is distributed among the bidders.

² Profit of buyer: $v_i - p + \frac{1}{3}p = v_i - b_1 + \frac{1}{3}b_1 = v_i - \frac{2}{3}b_1$

² Profit of non-buyers: $\frac{1}{3}p = \frac{1}{3}b_1$

Altogether you play 36 auctions. In each auction the bidder groups are formed randomly. After you have placed your bid you get information about the price, your private reselling value, whether or not you bought, and how much you have earned. Any decision you make is anonymous and can not be related to your person. If you have questions, please, don't ask loud, but raise your hand. We will then clarify problems privately.

Wiederverkaufswert	Gebot
50	<input type="text"/>
60	<input type="text"/>
70	<input type="text"/>
80	<input type="text"/>
90	<input type="text"/>
100	<input type="text"/>
110	<input type="text"/>
120	<input type="text"/>
130	<input type="text"/>
140	<input type="text"/>
150	<input type="text"/>

ERSTPREIS-AUKTION

Bitte geben Sie Ihre Gebotsfunktion ein:

OK

Figure B.1: Example of a bidding screen.

B. Sample screen shots

The screen in Figure B.1 was used by the subjects to place their bids. In the upper right corner one finds information about the type of the game. When subjects face a certain type of game for the first time all bid fields are empty. In later rounds subjects faced their strategy which they used in the last play of the same game type. They did not have to retype their strategy if they did not want to change it. Figure B.2 is the screen which subjects receive after an auction. It informs the participant whether (s)he became the buyer, the price, the individual value v_i^a , the own bid for this value, and the payoff resulting from all this events.

Ergebnis der Auktion ZWEITPREIS-
AUKTION

Sie haben gekauft.

Der Preis, den Sie zahlen mußten, war:

Ihr ausgewählter Wiederverkaufswert war:

Ihr Gebot für diesen Wiederverkaufswert war:

Ihr Gewinn beträgt:

Figure B.2: Example of an information screen.